

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problems Mailbox.**

THIS PAGE BLANK (USPTO)

EP 17114

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0702330 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
20.03.1996 Bulletin 1996/12

(51) Int Cl.⁶: G06T 15/00

(21) Application number: 95306309.6

(22) Date of filing: 08.09.1995

(84) Designated Contracting States:
DE FR GB

(30) Priority: 14.09.1994 US 306043

(71) Applicant: XEROX CORPORATION
Rochester New York 14644 (US)

(72) Inventors:

- Lamping, John O.
Los Altos, CA 94024 (US)
- Rao, Ramana B.
San Francisco, CA 94112 (US)

(74) Representative: Reynolds, Julian David et al
Rank Xerox Ltd
Patent Department
Parkway
Marlow Buckinghamshire SL7 1YL (GB)

(54) Layout of node-link structure in space with negative curvature

(57) Layout data indicate positions in a negatively curved layout space for nodes in a hierarchical branch of a node-link structure. The layout data indicate a parent position for parent nodes and, for children that share a parent node, child positions approximately along a circle in the layout space with the parent position approximately at the circle's center. Adjacent child positions are separated by approximately a base spacing. The radii of circles within the branch together approximate a function that increases slowly with number of child nodes such that the radii and spacings along circles are all approximately uniform within the branch. The layout data can be obtained (60) from data defining the node-link structure. The layout data can be used to perform mappings (64), each obtaining positions for a subset of the nodes. The layout data can be used to present (122: Fig.4) a first representation of the node-link structure on a display. In response to a user signal indicating (124: Fig. 4) a change from a first display position near a first feature to a second display position, a second representation can be presented (126: Fig 4) that is perceptible as a changed continuation of the first. The second representation includes, near the second display position, a second feature representing the same part of the node-link structure as the first feature. The second representation can be obtained by a transformation of the layout space, which can be a discrete approximation of a hyperbolic plane.




FIG. 2

A2
EP 0702330 A2

Description

The present invention relates to data processing, and more particularly to layout of node-link structures.

Fairchild, K.M., Poltrack, S.E., and Furnas, G.W., "SemNet: Three-Dimensional Graphic Representations of Large Knowledge Bases," in Guindon, R., Ed., *Cognitive Science and its Application for Human Computer Interaction*, Lawrence Erlbaum, Hillsdale, N.J., 1988, pp. 201-233, describe SemNet, a three-dimensional graphical interface.

As shown and described in relation to Fig. 1-1, SemNet represents a knowledge base as a directed graph in a three-dimensional space, with elements of knowledge represented as labeled rectangles connected by lines or arcs.

One aspect of the invention alleviates a basic problem in layout of node-link structures. The information in a node-link structure tends to grow exponentially as one traverses the structure. For example, if the node-link structure is a tree or other hierarchical structure with a root node and levels of nodes below the root, the number of nodes, and therefore the amount of information in the tree, tends to increase exponentially as one goes from the root through the levels below the root. For some applications, such as displaying a representation of a node-link structure, it is useful to lay out the node-link structure in a space so that each node has a unique position spaced well apart from other nodes. Exponential growth makes it difficult to lay out a node-link structure in a space.

Some conventional techniques lay out a node-link structure in a two-, three-, or other n-dimensional Euclidean space. For example, the technique of Fairchild et al. described above, positions nodes in a three-dimensional space, in part because arcs would intersect if nodes were positioned in a plane. Robertson et al., in US-A 5,295,243, describe techniques in which a hierarchical structure with rotating substructures is positioned in three-dimensional space.

Layout in Euclidean space can, for example, begin at a root node and traverse the structure, positioning each node in space. A layout in Euclidean space tends to include large variations in density because of exponential growth. Nodes in branches tend to be much denser than nodes nearer the root. Large variations in density cause problems in using the layout because a user of a browser based on such a layout has difficulty understanding relationships between nodes in regions of different densities.

The techniques of the invention are based on the discovery that a node-link structure with exponentially growing branches can be laid out without large variations in density in a space with negative curvature. A space has negative curvature if parallel lines diverge. Examples include n-dimensional hyperbolic spaces, with a two-dimensional hyperbolic space being referred to as a "hyperbolic plane." In a hyperbolic plane, for example, a

circle's area is proportional to an exponential function of its radius, making it easy to obtain a uniform layout for a node-link structure that has exponentially growing branches. A layout space with negative curvature is referred to as a "discrete approximation" to a hyperbolic plane or other continuous space with negative curvature if positions in the layout space are only specified with limited precision, so that every position in the layout space maps to a position in the hyperbolic plane or other continuous space with negative curvature but not vice versa.

Each of the techniques uses layout data indicating positions of parts of a node-link structure in a layout space with negative curvature. The layout data can, for example, include coordinates indicating positions in the layout space. The layout data indicate positions in the layout space for each of a set of nodes that form a branch. For each of a set of parent nodes, lower level nodes that share the parent node lie approximately along a circle in the space or along an arc of the circle, while the parent node is positioned approximately at the center of the circle. As used herein, the term "along a circle" includes an arrangement along an arc of the circle. The positions of adjacent lower level nodes along each circle are separated by approximately a base spacing. The radii of the circles for all the parent nodes in the set together approximate a radial spacing function that increases slowly with number of lower level nodes such that the radii and the spacings between adjacent lower level nodes along circles are all approximately uniform within the branch.

An approximately uniform layout of a branch can be very useful. For example, in displaying representations of a node-link structure that includes the branch, further layout is unnecessary--representations can be obtained by performing rigid transformations on the layout space to obtain transformed position data and by mapping from the transformed positions to a planar Euclidean region that can be presented on a display, such as a circular, disk-shaped region or any other appropriately shaped region.

The layout space can be a discrete approximation of a hyperbolic plane, for example, and the layout data can indicate, for each node, two coordinates that together indicate a position in the hyperbolic plane. The node-link structure can be an acyclic graph such as a tree.

One aspect of the invention relates to the display of representations of a node-link structure. In one technique, the layout data can be used to present a first representation of the node-link structure on a display. In response to a signal from a user indicating a change from a first display position to a second display position, the technique then presents a second representation that is perceptible as a moved continuation of the first representation. The first representation includes a first feature near the first display position, representing a first part of the node-link structure, and the second representation

includes a second feature representing the first part, but near the second display position.

More particularly, the present invention provides a method according to claim 1 of the appended claims.

Preferably, the layout data indicate, for each of the set of nodes, first and second coordinates: the first and second coordinates together indicating a position in the hyperbolic plane. Preferably, each of the first and second coordinates is between zero and one: the act of operating the processor to use the layout data to present a first representation of the node-link structure comprising: using the layout data to obtain mapped data that indicate, for each of a subset of the set of nodes, the first and second coordinates indicated for the node by the layout data; the first coordinate indicating position along a first axis within a planar unit disk and the second coordinate indicating position along a second axis within the planar unit disk, the second axis being perpendicular to the first axis.

Preferably, the node-link structure is an acyclic graph, and more preferably a tree. Preferably, the branch is an exponentially growing branch.

Preferably, the act of using the layout data to present a first representation of the node-link structure comprises: using the layout data to obtain mapped data indicating positions in a region of the display for a subset of the nodes; and using the mapped data to provide first representation data to the display, the first representation data causing presentation of the first representation in the region of the display: the first representation including, for each node in the subset, a node feature representing the node. Preferably, the act of using the layout data to obtain mapped data comprises performing projective mapping to a planar unit disk, or performing conformal mapping to a planar unit disk. Preferably, the layout data indicate, for each of the mapping's subset of nodes, first and second coordinates: the first and second coordinates together indicating a position in the space with negative curvature: each of the first and second coordinates being between zero and one: the act of using the layout data to obtain mapped data comprising: obtaining mapped data that indicate, for each of the subset of nodes, the first and second coordinates indicated for the node by the layout data; the first coordinate indicating position along a first axis within the planar unit disk and the second coordinate indicating position along a second axis within the planar unit disk, the second axis being perpendicular to the first axis.

Preferably, the region of the display has a perimeter: the act of using the layout data to obtain mapped data comprising, for one of the nodes that is a child node of one of the parent nodes and that has descendant nodes: using the layout data to map the parent node's position in the layout space to a first display position in the region of the display: using the layout data to map the child node's position in the layout space to a second display position in the region of the display; and if the second display position is within a minimum spacing from the perimeter of the region and, from the relation between the

first and second display positions, it is unlikely that any of the descendant nodes are further than the minimum spacing from the perimeter, excluding the child node and its descendant nodes from the subset of nodes. Preferably, the region of the display is circular.

The act of presenting the first representation may comprise: obtaining first transformation data indicating a first rigid transformation of the layout space that moves the positions indicated by the layout data to first transformed positions; and using the first transformation data to obtain first transformed position data: the first transformed position data indicating transformed positions in the layout space for parts of the node-link structure, the transformed positions having a relation to the positions indicated by the layout data equivalent to the first rigid transformation: the first transformed position data indicating a position for the first part of the node-link structure such that a feature representing the first part would be presented near the first display position: the act of presenting a second representation comprising: obtaining second transformation data indicating a second rigid transformation of the layout space that would move a first transformed position that maps to the first display position to a second transformed position that maps to the second display position: and using the second transformation data to obtain second transformed position data: the second transformed position data indicating transformed positions in the layout space for parts of the node-link structure, the transformed positions having a relation to the positions indicated by the first transformed position data equivalent to the second rigid transformation: the second transformed position data indicating a position for the first part of the node-link structure such that a feature representing the first part would be presented near the second display position. Preferably, the second transformation data indicate a second rigid transformation that preserves in the second representation an orientation of features in the first representation that represent parts of the node-link structure. Preferably, the first part is a parent node with at least one child node that is related to the parent node through one link: the first transformed position data indicating positions for the parent node and for its child nodes so that the child nodes together have a first orientation with respect to the parent node: the second transformation data indicating a second rigid transformation: the second transformed position data indicating positions for the parent node and for its child nodes so that the child nodes' positions together have a second orientation with respect to the second parent node's position: the first and second orientations being approximately equal. Preferably, the node-link structure is a tree with a root node: the first part being the root node.

Preferably, the first and second representations each include a region of greater spacings in which node features have nearest node spacings that are in general perceptually greater than in other regions: the first representation including in its region of greater spacings a first

parent node with at least one child node that is related to the first parent node through one link: the first transformed data indicating transformed positions for the first parent node and for its child nodes so that the child nodes together have a first orientation with respect to the first parent node in the first representation; the second representation including in its region of greater spacings a second parent node with at least one child node that is related to the second parent node through one link: the second transformation data indicating a second rigid transformation such that the second transformed position data indicate positions for the second parent node and for its child nodes so that the child nodes together have a second orientation with respect to the second parent node in the second representation; the first and second orientations being approximately equal.

The method may further comprise a series of iterations in response to the signal; each iteration comprising: presenting an intermediate representation of the node-link structure; the intermediate representation including features representing parts of the node-link structure, the features including an intermediate first part feature representing the first part; the intermediate first part feature having a position between the first display position and the second display position; the iterations being performed such that the intermediate first part features in the intermediate representations and the second feature in the second representation are all perceptible as moved continuations of the first feature in the first representation.

Preferably, the first part is a node.

The first representation can be presented by using the layout data to obtain mapped data. The mapped data can indicate positions in a region of the display for a subset of the nodes in which the first representation is presented, and the first representation includes a node feature for each node in the subset. If a node maps to a position that is within a minimum spacing from the display region's perimeter and there is no chance the node has descendants further from the perimeter than the minimum spacing, the node and its descendants can be excluded from the subset for which node features are presented.

The mapped data can be obtained by mapping to a circular display region such as a planar unit disk, for example. Either of two canonical mappings can be performed from the hyperbolic plane to a planar unit disk--the projective mapping, or Klein model, in which lines in the hyperbolic plane are taken to lines in the unit disk but angles are distorted, and the conformal mapping, or Poincaré model, in which angles are preserved but lines in the hyperbolic plane are distorted into arcs on the unit disk. If the Poincaré model is also used as the representation of points in the hyperbolic plane, then the Poincaré mapping to the unit disk is the identity operation on the representation, so that a pair of values indicating a position in the hyperbolic plane can be used directly as coordinates along orthogonal axes in the unit

disk.

The second representation can be presented by obtaining transformation data indicating a rigid transformation of the layout space and then using the transformation data to obtain transformed position data. The rigid transformation would move a first layout position that maps to the first display position to a second layout position that maps to the second display position. Therefore, the transformed position data indicate a position for

5 the first part of the node-link structure such that a feature representing the first part would be presented near the second display position.

A rigid transformation can be found that preserves or establishes the orientation of certain features after 15 mapping. In one variation, the rigid transformation preserves the orientation of children of some node with respect to that node; the parent node can, for example, be the root node. In another, the rigid transformation establishes parent-child orientation in a region of greater spacings in the representations.

To preserve a perception of continuity, the technique 20 can also present intermediate representations. The intermediate representations can include intermediate first part features which, together with the second feature in 25 the second representation, are all perceptible as moved continuations of the first feature in the first representation.

Another aspect of the invention provides techniques 30 for obtaining layout data as described above. One technique uses node-link data defining a node-link structure to obtain layout data as described above. The technique then uses the layout data to perform two or more different mappings. The mapped data obtained by each mapping indicate positions for a subset of the set of nodes. The 35 positions can, for example, be display positions or positions in another appropriate space.

Layout data can be obtained by first obtaining, for a 40 parent node, a parent position and a parent wedge in the layout space: the parent wedge has a vertex at the parent position and is a part of the layout space between two rays extending from the vertex. Then parent node data indicating the parent position and the parent wedge can be used to obtain child node data for each of the child 45 nodes of the parent node. A child node's data can indicate a child position approximately along an arc within the parent wedge, and a child wedge with a vertex at the child position and extending between rays from the vertex. Each child wedge is contained inside the parent wedge and does not overlap any other child wedge. 50 While the children of a root node can be spread over any arc up to 360°, the children of parent nodes below the root node must be spread over arcs smaller than 360°.

In one implementation, the rays of each child wedge 55 diverge from the vertex at approximately the same angle as the rays of the parent wedge; in this case, adjacent child positions can be equally separated along the arc by the base spacing. In another, the rays meet at angles that depend on how large an angle the descendants of

each child node would occupy were the descendant nodes of all children uniformly spaced.

Yet another aspect of the invention is based on the recognition of a problem that arises in implementing the technique of obtaining transformation data as described above when the layout data indicate positions in a discrete approximation of a hyperbolic plane. If the transformation is a pure translation, a problem with rotation arises. Translation performed with respect to a line in a hyperbolic plane rotates points not on the line with respect to a point on the line, which can be counterintuitive. Furthermore, a series of translations that form a closed loop, each preserving orientation along its line of translation, also causes a counterintuitive rotation.

This aspect of the invention provides techniques for alleviating the rotation problem. Each technique provides a transformation that includes a rotation component. One technique rotates so that children nodes in a region of greater spacings are presented in a canonical orientation, such as always to the right of their parent node. Another technique rotates to preserve orientation of a root node of the node-link structure.

The above techniques can be implemented in a machine with memory storing program data indicating instructions the machine's processor executes. More particularly, the invention further provides a machine according to claim 4 or 8 of the appended claims.

The above techniques can also be implemented in an article of manufacture for use in a machine. The article of manufacture can include a storage medium and data stored by the storage medium. The stored data can include data units positioned on the storage medium so that a storage medium access device can access the data units and provide the data units to the machine's processor in a sequence. The sequence of data units indicates instructions the processor executes. More particularly, the invention further provides an article of manufacture, according to claim 5 or 9 of the appended claims.

The above techniques can also be implemented in a method that operates a first machine to transfer instruction data indicating such instructions over a network to a second machine with a processor that executes the instructions.

For example, a machine's processor can use layout data indicating positions in a layout space with negative curvature as described above to provide first representation data to cause a display to present a first representation of a node-link structure. The processor can receive user signal data indicating a change from a first display position near a first feature in the first representation to a second display position, where the first feature represents a first part of the node-link structure. In response, the processor can automatically present a second representation in which a second part feature near the second display position represents the first part. The second representation can be perceptible as a moved continuation of the first.

A machine's processor can also use node-link data defining a node-link structure to obtain layout data indicating positions in a layout space with negative curvature as described above. The processor can use the layout

5 data to perform two or more different mappings that obtain mapped data, with the mapped data for each mapping indicating positions for a subset of the nodes in a region of the display.

More particularly, the invention further provides a
10 method of operating a machine, according to claim 3 of the appended claims.

Preferably, the user input circuitry comprises a pointer control device that controls a pointer presented on the display; the signal from the user including a button click
15 when the pointer is at the second display position. Preferably, the signal from the user further includes an action moving the pointer to the second display position.

The invention further provides a method of operating a machine, according to claim 6 of the appended claims.

20 The invention further provides a method of operating a machine, according to claim 7 of the appended claims.

Preferably, the layout space is a discrete approximation of a hyperbolic plane. Preferably, the layout data indicate, for each of the set of nodes, first and second coordinates; the first and second coordinates together indicating a position in the hyperbolic plane. Preferably, the node-link structure is an acyclic graph. Preferably, the node-link structure is a tree. Preferably, the branch is an exponentially growing branch. Preferably, the branch includes a parent node that is the top level node, the parent node having two or more child nodes at a first lower level; the act of operating the processor to use the node-link
25 data to obtain layout data comprising: obtaining parent node data indicating, for the parent node, a parent position and a parent wedge in the layout space; the parent wedge having a vertex at the parent position; the parent wedge being a part of the layout space between two rays extending from the vertex; and using the parent node data to obtain child node data for each of the child nodes:
30

35 each child node's child node data indicating, for the child node, a child position in the layout space; each of the child positions being approximately along an arc that is within the parent wedge, the arc being part of a circle that is centered at the vertex.

40 Preferably, the parent wedge has a wedge angle formed by the two rays at the vertex; the wedge angle being less than 360°; each child node's child node data indicating, for the child node, a child wedge; each child node's child wedge having a vertex at the node's child position, the child wedge being a part of the parent wedge between two rays extending from the vertex, the two rays forming an angle approximately equal to the wedge angle at the vertex; the child wedges not overlapping. Preferably, the branch includes descendant nodes
45

50 at lower levels below the first lower level, each descendant node being related to first or second ones of the children nodes through one or more links; more of the descendant nodes at one of the lower levels below the first

55

lower level being related to the first child node than to the second child node so that if the descendant nodes were uniformly spaced, the descendant nodes related to the first child node would occupy a larger angle than the descendant nodes related to the second children node: each child node's child node data indicating, for the child node, a child wedge: each child node's child wedge having a vertex at the node's child position, the child wedge being a part of the parent wedge between two rays extending from the child wedge's vertex, the two rays forming a wedge angle at the child wedge's vertex; the wedge angle of each child node's child wedge depending on the angle that would be occupied by descendant nodes related to the child node, the first child node having a larger wedge angle than the second child node: the child wedges of the child nodes not overlapping.

Preferably, the parent node of the branch is the root node of a tree: the two rays being identical so that the parent wedge includes all of the space with negative curvature.

Preferably, adjacent child positions are equally separated along the arc by the base spacing..

The invention further provides a method of operating a machine, according to claim 10 of the appended claims.

The techniques described above are advantageous because an exponentially growing node-link structure can be laid out uniformly in a space with negative curvature. Once it is laid out in a hyperbolic plane, for example, there is no need to lay it out again--the same layout can be transformed to produce other desired representations. The hyperbolic plane can be mapped to a display region such as a disk to show a part in greater detail and to show other parts with less detail, so that the context of the detailed part is visible. Furthermore, any transformation between two representations can be done in the same amount of time, allowing high quality animated navigation. A sequence of images presented at animation speeds can provide perception of continuity, while changing positions of parts of the structure. Therefore, a user can indicate a change in position of a part and a machine can automatically present an image in which the part is moved as indicated. This makes it possible for the user to change focus to any part of a structure.

Embodiments of the invention will now be described, by way of example, with reference to the accompanying drawings, in which:

Fig. 1 is a schematic flow diagram showing how node-link data can be used to obtain layout data indicating positions in a space with negative curvature, from which two or more different mappings can be performed:

Fig. 2 is a flow chart showing general acts in obtaining layout data and performing mappings as in Fig. 1:

5

Fig. 3 is a schematic flow diagram showing how first and second representations of a node-link structure can be presented, with the second representation automatically presented in response to a signal indicating a change from a first display position to a second display position in the first representation:

10

Fig. 4 is a flow chart showing general acts in presenting first and second representations as in Fig. 3:

15

Fig. 5 is a schematic diagram showing general components of a machine that obtains layout data and performs mappings as in Fig. 1 and presents first and second representations as in Fig. 3:

20

Fig. 6 is a schematic block diagram showing components of a system that can layout a node-link structure in a layout space with negative curvature, then map to present representations of the node-link structure on a display:

25

Fig. 7 is a flow chart showing acts by which the machine in Fig. 6 can provide a browser interface by transforming and mapping positions for a subset of nodes not too near an edge in response to requests for position changes:

30

Fig. 8 is a flow chart showing acts by which the machine in Fig. 6 can provide a browser interface by successively transforming and mapping positions for all nodes in response to requests for position changes:

35

Fig. 9 is a flow chart showing acts in a layout operation by the system of Fig. 6 implementing the technique of Fig. 7:

40

Fig. 10 is a flow chart showing acts that can obtain a node's distance to children; a child's position, wedge, and room bound; and a node's radius in Fig. 9:

45

Fig. 11 is a flow chart showing acts in mapping from layout to display by the system of Fig. 6 implementing the technique of Fig. 7:

50

Fig. 12 is a flow chart showing acts that can obtain a node's mapped position in a unit disk and present a link feature in Fig. 11:

55

Fig. 13 is a flow chart showing acts that can present a node feature in Fig. 11:

60

Fig. 14 is a flow chart showing acts in responding to a signal requesting change of position by the system of Fig. 5 implementing the technique of Fig. 7:

65

Fig. 15 is a flow chart showing acts that can obtain

a translator that preserves orientation of a root node in Fig. 14; and

Fig. 16 is a flow chart showing acts that can obtain a translator that orients children of a node at the center of a region of greater spacings the right in Fig. 14.

A "graph" is a node-link structure in which each link relates two nodes. An "acyclic graph" is a graph in which the links provide only one path between any two nodes in the graph. A "directed graph" is a graph in which each link indicates direction between the nodes it relates, with one node being a source of the link and the other being a destination. A "tree" is an acyclic directed graph with exactly one root node such that every other node in the tree can be reached by only one path that begins at the root node and follows each link in the path in its indicated direction.

A "branch" of a node-link structure is a set of nodes that forms a tree within the node-link structure if the links are treated as relating pairs of nodes and as indicating direction. A branch therefore includes two or more levels, with the "top level node" being the node that is the root node of the tree and "lower level nodes" being nodes at one or more levels of the tree below the top level node. Each lower level node has a "parent node" at the next higher level to which the lower level node is related through one link. A parent node has a set of "child nodes" at the next lower level to each of which the parent node is related through one link. The child nodes of a parent "share" the parent node. A parent node also has a set of "descendant nodes" that includes its child nodes and all of their child nodes, and so forth.

An "exponentially growing branch" of a node-link structure is a branch in which the number of nodes at successively lower levels increases approximately as an exponential function of the number of links m relating nodes at a lower level to the top level node of the branch. For example, if parent nodes within the branch have an average of three children, the number of nodes at the m th lower level would be approximately 3^m .

A "representation of a node-link structure" is a graphical representation that represents the node-link structure. In a representation of a node-link structure, for example, link features can be lines, such as arcs or straight lines, that extend between node features.

A "bounded node feature" is a node feature that has a perceptible boundary. The "center of area" of a bounded node feature is the center of area of the region within the node feature's boundary. The position of a bounded node feature's center of area can therefore be computed from the node feature's boundary or estimated by viewing the representation.

The "nearest other node feature" of a first bounded node feature in a representation is a second bounded node feature whose center of area is spaced from the first node feature's center of area by a distance no great-

er than the spacing from the first node feature's center of area to any other bounded node feature's center of area. The distance is referred to herein as the node feature's "nearest node spacing." A bounded node feature 5 may have more than one nearest other node feature, all with centers of area at the nearest node spacing.

Bounded node features in a first region have nearest node spacings that are "in general perceptibly greater" than in a second region if a viewer can see that the nearest 10 node spacings are generally greater in the first region than in the second. A region in which bounded node features have nearest node spacings that are in general perceptibly greater than in other regions of a representation may be referred to as a "region of greater spacings."

In a representation of a node-link structure, a relation 15 between a parent node feature's display position and a child node feature's display position is such that "it is unlikely that any of the descendant nodes are further than" a minimum spacing from the perimeter of a region if, given display positions of a parent node feature and a child node feature and a way in which the display positions were obtained, it is unlikely that a descendant of the child would be positioned further than the minimum spacing from the perimeter. For example, in an implementation described below, a test for unlikelihood would be whether the relation between the two display positions is such that the child node feature's display position is nearly as close to the perimeter as the parent node feature's display position.

A "space with negative curvature" is a space in 20 which parallel lines diverge. Therefore, through any position in a space with negative curvature that is not on a given straight line, there are multiple other straight lines parallel to the given straight line. An example of a space with negative curvature is hyperbolic n -space. A "hyperbolic plane" is a hyperbolic 2-space.

A "discrete approximation" of a continuous space is 25 a space in which positions can be uniquely specified by a subset of the continuous range of numbers for the continuous space, and the subset is such that positions in the discrete approximation can only be specified with limited precision in comparison with the continuous space. For example, if positions in a continuous space can be specified by the real numbers, positions in a discrete approximation might be specified only by the integers or by n -bit numbers, and so forth.

An item of data indicates positions "approximately 30 along a circle" or "approximately along an arc" in a layout space if an arc of an ellipse or a circle can exist within the layout space such that each position is closer to the arc than to an adjacent one of the positions. Adjacent positions that are approximately along a circle are "separated by approximately a base spacing" if the distances in the layout space between the adjacent positions are 35 within an order of magnitude of the base spacing.

An item of data indicates a position "approximately 40 at the center" of a circle in a layout space if the position is closer to the center of the circle than to the circle.

Within a branch that includes two or more parent nodes for each of which an item of data indicates positions for child nodes sharing the parent node approximately along a circle and a position for the parent node approximately at the center of the circle, the radii of circles and spacings between adjacent child positions along circles are "all approximately uniform within the branch" if the greatest and the smallest of the radii and spacings are within an order of magnitude of each other. The circles have radii that, together, "approximate a function that increases slowly with number of lower level nodes" if mean lengths of radii of the circles that have each number of child nodes form a curve that approximates a function with monotonic positive slope and whose slope decreases with increasing number of children.

Figs. 1-5 show general features of the invention. Fig. 1 shows how node-link data defining a node-link structure can be used to obtain layout data indicating positions in a layout space with negative curvature, from which two or more different mappings can be performed. Fig. 2 shows general acts in using node-link data to obtain layout data and perform mappings as in Fig. 1. Fig. 3 shows how first and second representations can be presented, with the second representation being automatically presented in response to signals from a user indicating a change from a first display position in the first representation to a second display position. Fig. 4 shows general acts in presenting first and second representations as in Fig. 3. Fig. 5 shows general components of a machine that can obtain layout data and use it to perform mappings as in Fig. 1 and can present first and second representations as in Fig. 3.

In Fig. 1, node-link data 10 define node-link structure 20, part of which is illustratively shown. As shown, node-link structure 20 includes nodes and links that relate the nodes. Fig. 1 shows an exponentially growing branch of node-link structure 20, with a top level node 22 related to each of lower level nodes 24, 26, and 28 through one link. As suggested by the dashed lines, each of lower level nodes 24, 26, and 28 is in turn related to three nodes at a next lower level. If this rate of increase continues, the number of nodes at a lower level of the exponentially growing branch would be approximately 3^m , where m is the number of links relating nodes at the level to top level node 22. As suggested in Fig. 1, it is difficult to lay out an exponentially growing branch on a two-dimensional Euclidean plane because of space constraints.

Layout data 30, however, indicate positions in negatively curved space 32 for nodes in the exponentially growing branch of node-link structure 20. Space 32 is schematically illustrated with position 34 of top level node 22 at the vertex of a wedge that extends between rays 36 and 38 from the vertex. This top level wedge is allocated to top level node 22. Lower level nodes 24, 26, and 28 can have positions 42, 44, and 46 approximately along circle 48 whose center in space 32 is approximate-

ly at position 34, the vertex of the top level wedge.

Positions 42, 44, and 46 all lie within the top level wedge and all the other descendants of top level node 22 at lower levels can also be positioned within the top level wedge. For example, nodes at a next lower level of the exponentially growing branch can similarly be approximately positioned along circles whose centers in space 32 are approximately at positions 42, 44, and 46. Because space 32 is negatively curved, each of lower level nodes 24, 26, and 28 can be allocated a lower level wedge between rays that diverge from the vertex of the wedge at approximately the same angle as the angle at which rays 36 and 38 diverge from position 34, even though the lower level wedges do not overlap each other and all are contained within the top level wedge.

Adjacent child positions 42 and 44 on circle 48 are separated by distance d_1 , while adjacent child positions 44 and 46 are separated by distance d_2 . Each of the distances d_1 or d_2 is approximately equal to a base spacing, and radius d_0 of circle 48 and the radii of other circles in the branch together approximate a function that increases slowly with number of lower level nodes such that the radii and spacings between adjacent child positions along circles are all approximately uniform within the branch.

As shown in Fig. 1, layout data 30 can be used to perform two or more different mappings. Each mapping obtains mapped data, as illustrated by zeroth mapped data 50 and Nth mapped data 52. The mapped data obtained by each mapping indicate positions for a subset of the nodes in node-link structure 20. The positions could, for example, be display positions or positions in another space other than the layout space.

In Fig. 2, the act in box 60 begins by obtaining node-link data defining a node-link structure. The node-link structure includes nodes and links, with each link relating at least two nodes.

The act in box 62 uses the node-link data obtained in box 60 to obtain layout data indicating, for nodes in a branch, positions in a layout space with negative curvature. For each of a set of parent nodes, the positions of children sharing a parent node are approximately along a circle centered approximately at the position of the parent node, with adjacent children nodes separated by distances that approximate a base spacing. The radii of circles in the branch together approximate a function that increases slowly with number of children nodes sharing a parent such that the radii and spacings between adjacent child positions along the circles are all approximately uniform within the branch.

The act in box 64 then uses the layout data from box 62 to obtain mapped data for at least two different mappings. The mapped data for each mapping indicate positions for a subset of the nodes.

In Fig. 3, layout data 70 indicate positions in a layout space with negative curvature for parts of a node-link structure. Layout data 70 is used to present image 72 with first representation 74 of the node-link structure. Im-

age 72 illustratively also includes pointer 76, which can be controlled by a mouse or other pointer control device to provide signals indicating display positions.

As suggested by the dashed line in Fig. 3, an uncertain amount of time passes before the user provides signal 78 indicating a change from a first display position to a second display position. First representation 82 has a first feature representing a first part of the node-link structure near the first display position. For example, one of node features 80, 82, and 84 could be near the first display position.

In response to signal 78, image 86 is presented with second representation 88 of the node-link structure. Second representation 88 is perceptible as a changed continuation of first representation 74.

In second representation 88, a second feature representing the first part of the node-link structure is near the second display position. For example, if node feature 80 in first representation 82 was near the first display position, and node feature 90 represents the same node, the second display position could have been indicated by a dragging signal from the position of node feature 80 to the position of node feature 90. Similarly, if node features 92 and 94 represent the same nodes, respectively, as node features 82 and 84, the second position could have been indicated by a dragging signal from the position of node feature 82 or node feature 84 to the position of node feature 92 or node feature 94, respectively. Also, the second position could have been indicated by clicking on node feature 82 in first representation 74 to indicate a second position at the center of a region of greater spacings, as suggested by the increase in size from node feature 82 to node feature 92.

Fig. 4 shows general acts implementing the operations illustrated in Fig. 3. In the act in box 120, layout data is obtained, indicating positions in a layout space with negative curvature for parts of a node-link structure.

The act in box 122 uses the layout data from box 120 to present a first representation of the node-link structure. The first representation includes a first feature representing a first part of the node-link structure, and the first feature is near a first display position.

After an uncertain amount of time, as indicated by the dashed line, the act in box 124 receives a signal from a user indicating a change from the first display position to a second display position.

The act in box 126 responds to the signal in box 124 by presenting a second representation of the node-link structure. The second representation includes a second feature representing the first part of the node-link structure. The second feature is near the second display position.

Machine 150 in Fig. 5 includes processor 152 connected for receiving data indicating user actions from user input circuitry 154 and for providing data defining images to display 156. Processor 152 is also connected for accessing node-link data 158, which define a node-link structure as described above in relation to Fig.

2. Processor 152 is also connected for receiving instruction data 160 indicating instructions through instruction input circuitry 162, which can illustratively provide instructions received from connections to memory 164, storage medium access device 166, or network 168.

In executing the instructions indicated by instruction data 160, processor 152 uses node-link data 158 to obtain layout data indicating positions of parts of the node-link structure in a layout space with negative curvature, as in box 62 in Fig. 2 and box 120 in Fig. 4. Processor 152 then uses the layout data to present a first representation of the node-link structure on display 156 as in box 122 in Fig. 4, which can be done by performing a first mapping for a subset of the nodes to a region of the display, as in box 64 in Fig. 2. When processor 152 receives data indicating a user signal from user input device 154 and the signal indicates a change from a first display position to a second display position as in box 124 in Fig. 4, processor 152 can automatically present a second representation of the node-link structure on display 106, with the position of a feature near the first display position changed to near the second display position, as in box 126 in Fig. 4. The second representation can be presented by performing a second mapping as in box 64 in Fig. 2.

As noted above, Fig. 5 illustrates three possible sources from which instruction input circuitry 162 could receive data indicating instructions--memory 164, storage medium access device 166, and network 168.

Memory 164 could be any conventional memory within machine 150, including random access memory (RAM) or read-only memory (ROM), or could be a peripheral or remote memory device of any kind.

Storage medium access device 166 could be a drive or other appropriate device or circuitry for accessing storage medium 170, which could, for example, be a magnetic medium such as a set of one or more tapes, diskettes, or floppy disks; an optical medium such as a set of one or more CD-ROMs; or any other appropriate medium for storing data. Storage medium 170 could be a part of machine 150, a part of a server or other peripheral or remote memory device, or a part of a software product. In each of these cases, storage medium 170 is an article of manufacture that can be used in machine 150. Data units can be positioned on storage medium 170 so that storage medium access device 166 can access the data units and provide them in a sequence to processor 152 through instruction input circuitry 162. When provided in the sequence, the data units form instruction data 160, indicating instructions as illustrated.

Network 168 can provide instruction data 160 received from machine 180. Processor 182 in machine 180 can establish a connection with processor 152 over network 168 through network connection circuitry 184 and instruction input circuitry 162. Either processor could initiate the connection, and the connection could be established by any appropriate protocol. Then processor 182 can access instruction data stored in memory 186 and

transfer the instruction data over network 168 to processor 152 so that processor 152 can receive instruction data 160 from network 168. Instruction data 160 can then be stored in memory 164 or elsewhere by processor 152, and can be executed.

The general features described above could be implemented in numerous ways on various machines to obtain and use layout data indicating positions in a layout space with negative curvature for parts of node-link structures. An implementation described below has been implemented on a Sun SPARCStation running a Unix/X operating system. Another implementation has been done on an Apple Macintosh personal computer. The Unix implementation executes instructions written in a Common Lisp programming environment extended with the Common Lisp Interface Manager (CLIM). Commercial Common Lisp environments and CLIM extensions are available from Lucid Inc. and Franz Inc.

Fig. 6 illustrates components of systems in which the invention has been implemented. Machine 200 in Fig. 6 includes central processing unit (CPU) 202, a microprocessor or other appropriate processor. CPU 202 is connected to receive data indicating user signals from user input devices that illustratively include keyboard 204 and mouse 206. CPU 202 is also connected to provide data to cause presentation of images on display 208. Images presented on display 208 can include pointer 210 controlled by mouse 206 or another pointer control device.

CPU 202 is connected to one or more memory components that together provide memory environment 220. The memory components conventionally include resident RAM connected directly to CPU 202, one or more internal disk drives or other storage medium access devices connected to CPU 202 through an internal bus, and one or more external memory components such as peripheral devices or remote servers that include storage medium access devices and that are connected to CPU 202 through external connections of machine 200. CPU 202 conventionally accesses data stored in any memory component in environment 220 by providing addresses.

Fig. 6 illustratively shows program components and data components stored in memory environment 220; program and data components are illustratively shown separately, but need not be stored separately within a given memory component. Program components include data that indicate instructions CPU 202 can execute, and can be stored in object code form or source code form; if appropriate, the program components can also include any compilers and interpreters needed to obtain instructions from program components stored in source code form. In executing instructions indicated by program components, CPU 202 accesses and uses data components.

In executing instructions indicated by operating system program 230, CPU 202 performs operations that transfer data between CPU 202 and other components of machine 200. Such operations can include, for example, mapping logical addresses in the memory space of

CPU 202 to physical addresses for accessing a memory component in environment 220.

In executing instructions indicated by Common Lisp environment program 232, which includes a CLIM library 5 and user interface extensions, CPU 202 provides the standard CLIM interface defined in the CLIM 2.0 specification available from Franz Inc. As a result, CPU 202 can execute CLIM source code, so that browser interface program 234 and hyperbolic geometry program 236 can be implemented in Common Lisp extended with CLIM.

In executing instructions indicated by browser interface program 234, CPU 202 first initializes machine 200. Then CPU 202 uses node-link data 240 defining a node-link structure to obtain layout data 242, indicating 15 positions for parts of the node-link structure in a hyperbolic plane. Then CPU 202 uses layout data 242 to obtain mapped data 244 for an initial representation of the node-link structure. CPU 202 provides data to display 208 to cause presentation of the initial representation.

20 In executing instructions indicated by browser interface program 234, CPU 202 also receives data indicating user signal events from keyboard 204 and mouse 206. CPU 202 sets up, maintains, and handles a queue of steps for responding to an event by accessing event 25 queue data 246. When the steps responding to a user signal event are completed, CPU 202 can set up steps for responding to the next event.

In executing instructions indicated by browser interface program 234, CPU 202 can also modify animation 30 parameters by accessing animation parameter data 248. Or CPU 202 can use layout data 242 to obtain mapped data 244 for a changed representation of the node-link structure, then provide data to display 208 to cause its presentation.

35 CPU 202 can execute instructions indicated by hyperbolic geometry program 238 for various purposes, such as to obtain layout data 242, to obtain mapping data 244, and to map from a position of pointer 210 to a part of the node-link structure. As shown in Fig. 6, a representation can include a background 250 on which a figure 252 is shown, and a position of pointer 210 within figure 252 can be mapped to the nearest node, for example. The operations of CPU 202 described above can be implemented in many ways to provide a browser for 40 a node-link structure laid out in a space with negative curvature. Fig. 7 shows how the machine in Fig. 6 can provide a browser by transforming positions for nodes not too near an edge. Fig. 8 shows how the machine in Fig. 6 can provide a browser by successively transforming 45 layout data for all nodes.

The general approach in Figs. 7 and 8 is that positions of nodes can be transformed in a layout space; then, a representation can be presented by mapping from the layout space to display positions. As mentioned 50 above, the layout space can be a hyperbolic plane or another appropriate space with negative curvature. The display positions can be positions within a circular display region, to facilitate mapping between the hyperbolic

plane and the display region.

The act in box 260 in Fig. 7 obtains layout data indicating a position in a layout space for each node in a node-link structure. The layout space can, for example, be a hyperbolic plane. The act in box 260 also initializes a current transformation that can be performed on the layout data to obtain transformed positions. The initial current transformation could, for example, be a null transformation that does not change the positions of the nodes in the layout space.

The act in box 262 then performs the current transformation on the layout data from box 260 to obtain transformed data indicating transformed positions for each node except those nodes with positions in the layout space that are too near an edge. The edge bounds the part of the layout space that is presented when a representation of the node-link structure is displayed.

The act in box 264 then maps the transformed positions from box 262 to obtain a position on a display for each node that is not too near the edge. The act in box 264 then presents a representation on the display. The representation includes, for each node that is not too near the edge, a node feature at the node's display position.

As indicated by the dashed line in Fig. 7, the act in box 264 can be followed after an unspecified time by the act in box 270, which receives a request from a user for a position change. The request includes an indication of at least one position on the display.

The act in box 272 maps the display position(s) indicated in box 270 back to the layout space to obtain start and end positions in layout space. For example, if only one display position is indicated, as by a mouse click, the position of the pointer when the click occurs can be used to obtain the start position and the end position can be the position in layout space of the center of the representation. Or, if two display positions are indicated, as during a dragging motion from a button press to a button release, the position of the pointer when the button press occurs can be used to obtain the start position and the position when the button release occurs can be used to obtain the end position.

The act in box 274 obtains a transformation of the layout space that moves the start position to the end position. The act in box 276 combines the current transformation with the new transformation from box 274 to obtain a combined transformation. The combined transformation becomes the current transformation when the acts in boxes 262 and 264 are performed to present another representation.

As described below, the technique in Fig. 7 has been implemented to present a sequence of representations in response to a click rather than just one representation. This modification can provide a perception of continuous motion from the start position to the end position.

The technique in Fig. 8 is similar to Fig. 7, with the act in box 280, for example, being the same as the act in box 260. A key difference, however, is that the act in

box 282 obtains a transformed position for each node in the node-link structure, including those that are treated as too near the edge in box 262. As a result, the technique of Fig. 8 is not as efficient as the technique of Fig. 7 for large node-link structures.

The act in box 284 maps the transformed positions from box 282 to display positions, as in box 264 in Fig. 7, and uses the display positions to present a representation as in box 264.

- 10 The act in box 290 receives a request as in box 270. The act in box 292 responds by first updating the layout data to be the transformed data from box 282 rather than the layout data from box 280. Then the act in box 294 maps from the display position(s) indicated in box 290 to the transformed layout space to obtain start and end positions. The act in box 294 obtains a new transformation that moves the start position to the end position, and updates the current transformation to be the new transformation before the acts in boxes 282 and 284 are performed to present another representation.

As noted above, the technique of Fig. 7 can handle large node-link structures better than the technique of Fig. 8 because the act in box 262 is not performed for nodes that are too near the edge. The technique of Fig. 25 has been found to be impractical for large structures. In addition, because some information is lost in obtaining a transformation, the technique of Fig. 7 is more precise, because it continues to use the original layout data from box 260.

30 The following implementation of CPU 202 operations follows the browser technique of Fig. 7, but the basic operations described below could also be used to implement the browser technique of Fig. 8.

The implementation described below lays out a 35 node-link structure in a uniform way in the hyperbolic plane and maps this plane onto a circular display region. The hyperbolic plane is a mathematical construct in which parallel lines diverge away from each other. This leads to the convenient property that the circumference 40 of a circle grows exponentially with its radius, which means that exponentially more space is available with increasing distance. Thus an exponentially growing node-link structure can be laid out in hyperbolic space in a uniform way, so that the distance between parents, 45 children, and siblings is approximately the same everywhere in the hierarchy.

While the hyperbolic plane is a mathematical object, it can be mapped in a natural way onto the unit disk, which provides a way to display it on a two-dimensional 50 (Euclidean) display. This mapping displays portions of the plane near the origin using more display space than other portions of the plane. Very remote parts of the hyperbolic plane get minuscule amounts of space near the edge of the disk. The projection onto the unit disk thus 55 provides a natural mechanism for assigning more space to a portion of the node-link structure while still embedding it in the context of the entire structure.

Translating the structure on the hyperbolic plane

provides a mechanism for controlling which portion of the structure receives the most space without compromising the illusion of viewing the entire structure. The structure is initially displayed with its root at the center of a region of greater spacings, but the implementation includes effective procedures for manipulating focus using pointer clicks as well as interactive dragging and for smoothly animating transitions across such manipulations. The context always includes several generations of parents, siblings, and children, making it easier for the user to explore the structure without getting lost.

The browser described below supports effective interaction with much larger structures than conventional hierarchy viewers and compares favorably with other novel browsers such as those described in Robertson et al., US-A 5,295,243. In a 600 pixel by 600 pixel window, a standard two-dimensional hierarchy browser can typically display 100 nodes, each with a three character text string. The browser described below can display 1000 nodes of which about the 50 nearest the focus can show from three to dozens of characters of text. Thus the can display up to ten times as many nodes while providing more effective navigation around a node-link structure. The scale advantage is obtained by the dynamic distortion of the structure display according to the varying interest levels of its parts.

Figs. 9 and 10 illustrate acts in layout operations. Fig. 9 illustrates how CPU 202 lays out a node-link structure in a hyperbolic plane: the technique of Fig. 9 can be used to implement the acts in box 260 in Fig. 7 and box 280 in Fig. 8. Fig. 10 illustrates how a distance to a node's children: a child's position, wedge, and room bound; and a node's radius can be obtained in Fig. 9.

The layout technique in Fig. 9 recursively lays out each node in a node-link structure based on local information. A node is allocated a wedge of the hyperbolic plane, angling out from itself, in which its descendants can be positioned. All of a node's children are positioned along an arc in that wedge, at an equal distance from the node, and far enough out so that the children are some minimum distance apart from each other. Each of the children then gets a sub-wedge for its descendants.

Because of the way parallel lines diverge in hyperbolic geometry, each child will typically get a wedge that spans about as big an angle as does its parent's wedge. To obtain a more compact layout for non-uniform structures, this is modified so that siblings that themselves have many descendants get more room than siblings that don't: this modification tends to put grandchildren more nearly the same distance from their grandparent.

The technique of Fig. 9 can also be modified by changing the angle occupied by the children of the root node. The children of the root node can be spread out over an entire 360° circle, or can all be put on one side. In some cases, the second option is better.

The act in box 350 in Fig. 9 begins by receiving a call to perform a layout operation with the handle of a root node in the node-link data, with coordinates of the

origin of a hyperbolic plane, with a wedge of the hyperbolic plane that is available for the descendants of the root node, and with a room bound. The layout operation of Fig. 9 can be used with node-link data defining an acyclic hierarchical node-link structure such that the handle of the node-link data is also the root node's handle. The wedge of the hyperbolic plane can be defined by a wedge angle and a direction. For a radial orientation, the wedge angle can be 360° and the direction does not matter, but

for a canonical orientation such as rightward, the wedge angle is smaller, such as 120°, and the direction determines the orientation. The room bound can be initialized to 0.9 for the root node.

The act in box 352 tests whether the node currently being handled has children. If so, the act in box 354 obtains a distance to the children in the hyperbolic plane, as described below. Then the act in box 360 begins an iterative loop that handles each of the current node's children in turn.

The act in box 362 begins each iteration by obtaining, for the next child node, a position in the hyperbolic plane, a wedge, and a room bound, as described below. The act in box 364 then makes a recursive call for a layout operation with the handle of the next child and with its position, wedge, and room bound from box 362. The act in box 366 returns from the recursive call and the act in box 368 then adds the child's handle to a list of children for the current node.

When all of the current node's children have been handled, the act in box 370 obtains a radius for the current node, by taking either the room bound received in the call in box 350 or box 364 or one-half the distance to the current node's children from box 354, whichever is smaller, as described below in relation to Fig. 10. If the current node has no children, the act in box 370 takes the room bound from box 350 as the radius.

The act in box 372 then creates a data structure for the current node. The current node's data structure includes its position and its radius from box 370, and, if the current node has children, its node data structure also includes a handle or other link to the list of its children from box 368. The data structure can also include other relevant information. Then, the act in box 374 ends by returning the handle of the current node's data structure.

Fig. 10 shows how the acts in boxes 354, 362, and 370 in Fig. 9 can be implemented.

In the following description, each distance between a pair of points is encoded as the distance between the points in the Poincaré model with one point at the origin. This distance is the inverse hyperbolic tangent of the distance between the points in hyperbolic space.

The act in box 390 begins by obtaining the distance to a node's children by obtaining a value for space needed, using a base spacing S that is the same for all nodes. The space needed depends on the angles that would be occupied by the descendants of each of a node's children if the descendants were uniformly spaced: these angles are referred to as "weights" herein. A value for

each child's weight can be obtained by taking the larger of 1.0 and $\ln(1.0 + \text{Grandchildren})$, where the variable Grandchildren is equal to the summation of the greater of 1.0 and $\ln(1.0 + (\text{Number of Children}))$ for all grandchildren of the node that are children of the child.

To obtain the value for space needed, the act in box 390 also obtains an angle fraction by dividing the total angle available by the sum of the weights for all the children of the current node. The angle fraction is multiplied by the smallest of the weights to obtain an angle unit α . The space needed is then equal to $((\text{ssin}^2 + 1.0)^{\frac{1}{2}} \cdot \text{ssin})$, where $\text{ssin} = (((1.0 - S^2)/2S) \cdot \sin \alpha)$.

The act in box 392 then determines whether the space needed from box 390 is larger than the minimum space, which is the hyperbolic sum of the base spacing S with itself, calculated as $2S/(1 + S^2)$. If so, the act in box 394 sets the distance equal to the space needed, but if not, the act in box 396 sets the distance equal to the minimum space.

The act in box 398 then initializes a value for previous width W_o to zero: a value for down room, which can be initialized to the hyperbolic half of the distance D from box 394 or 396, calculated as $D/(1.0 + (1.0 - D^2)^{\frac{1}{2}})$; and values indicating direction, which can be initialized using values indicating the angular width W of the current node's wedge and the direction of the current node's wedge in the form of (θ_x, θ_y) , x- and y-coordinates of a point on the unit disk. A starting direction (θ_x, θ_y) for calculating the first child's direction can have the calculated values $((\theta_x(\cos(-W))) - (\theta_y(\sin(-W)))), ((\theta_y(\cos(-W))) + (\theta_x(\sin(-W))))$.

Then the act in box 360 begins an iterative loop as described above in relation to Fig. 9. Part of the iterative loop is shown in boxes 400, 402, 404, and 406.

The act in box 400 obtains preliminary width and direction of the next child node's wedge, then uses the preliminary width to update last width. The act in box 400 also obtains the next child node's position, wedge, and room available.

The act in box 400 can obtain preliminary width angle W_n as the product of the child's weight and the angle fraction, both obtained in box 390. The act in box 400 can use (θ_x, θ_y) and the previous width angle W_o , either from box 398 or from the previous iteration, to calculate child direction (θ_x, θ_y) having coordinate values equal to $((\theta_x(\cos(W_o + W_n))) - (\theta_y(\sin(W_o + W_n)))), ((\theta_y(\cos(W_o + W_n))) + (\theta_x(\sin(W_o + W_n))))$. Then the act in box 400 updates previous width W_o to have the value of the preliminary width W_n .

The act in box 400 can obtain the next child node's position using the position (x, y) of the parent node, from box 350 if the parent node is the root node or from box 364 when performed for the parent node: using the child direction (θ_x, θ_y) ; and using the distance D from box 394 or 396. The position include values $(x, y, \theta_x, \theta_y)$ that are calculated using transformations.

A rigid transformation of the hyperbolic plane represented by a Poincaré model is determined by a point p

and an angle α , where p indicates the final position of the point that was the origin prior to the transformation and where α represents a rotation. The transformation of a point (x, y) is then $((\theta z + p)/(P^* \theta z + 1))$, where z is

- 5 a complex number with real part x and imaginary part y : P is the complex number representing p : θ is the complex number of magnitude one representing rotation by the angle α ; and P^* is the complex conjugate of P . The real and imaginary parts of the result are the x - and y -coordinates, respectively, of the transformed point p in the Poincaré model.

Given values defining a point p and the cosine and sine of an angle α , together represented by the four coordinates $(x, y, \theta_x, \theta_y)$, a transformation can be obtained

- 15 by obtaining P as a complex number with real part x and imaginary part y and θ as a complex number with real part θ_x and imaginary part θ_y .

To obtain a transformation from the origin to a point (x, y) , a transformation can therefore be obtained from

- 20 the point and direction $(x, y, 1.0, 0.0)$, where the coordinates 1.0 and 0.0 indicate the rightward direction, the zero direction from which other directions are measured. This transformation is referred to herein as the "translator from the origin." Conversely, to obtain a transformation
- 25 from a point (x, y) to the origin, a transformation can therefore be obtained from the point and direction $(-x, -y, 1.0, 0.0)$. This transformation is referred to herein as the "translator to the origin."

The act in box 400 can obtain a first translator from

- 30 the origin to the position of the parent node, then obtain the new x - and y -coordinates by applying the first translator to the x - and y -components of the distance D , $(D\theta_x, D\theta_y)$, where (θ_x, θ_y) is the child direction described above. The act in box 400 can then obtain a second
- 35 translator to the origin for the new x - and y -coordinates, and apply the second translator to the results of applying the first translator to the child direction (θ_x, θ_y) , to obtain the new values for (θ_x, θ_y) .

The new values for (θ_x, θ_y) indicate the direction for

- 40 the child node's wedge. The act in box 400 can obtain the angle of the next child's wedge by using the preliminary width W_n and the distance D from box 394 or box 396. The act in box 400 can obtain a translator to the origin for the coordinates $(D, 0.0)$, then apply this translator to $\cos W_n$ and $\sin W_n$ to obtain values that have a ratio R . The angle of the wedge can then be calculated as $((\pi/2) - \arctan R)$.

Finally, the act in box 400 can obtain the room available to be used in obtaining the room bound to be passed

- 50 in the recursive call in box 364. The room available can be calculated using the distance D from box 394 or box 396 and the preliminary width W_n . The room available can then be calculated as $((s^2 + 1.0)^{\frac{1}{2}} - s)$, where $s = ((1.0 - D^2)/2D \cdot \sin W_n))$.

The act in box 402 then branches based on whether the room available from box 400 is greater than the down room from box 398. If so, the act in box 404 sets room bound to the value of room available. If not, the act in

box 406 sets room bound to the value of down room.

Then, when every child of the current node has been handled by an iterative loop, the act in box 410 branches based on whether the room bound received in box 350 is greater than the down room from box 398. If so, the act in box 412 sets the node's radius to the value of room bound. If not, the act in box 414 sets the radius to the value of down room.

The operations in Fig. 10 result in approximately uniform spacings between nodes within a branch, because the base spacing S is a minimum below which spacings do not occur and because other spacings turn out to be not much larger than S , because they grow only slowly with the number of children of a node.

Figs. 11-13 illustrate operations mapping from a hyperbolic plane to a display region. Fig. 11 shows how mapping from layout space to display positions can be performed; the technique of Fig. 11 can be used to implement the acts in boxes 262 and 264 in Fig. 7, and a similar technique could be used to implement the acts in boxes 282 and 284 in Fig. 8. Fig. 12 shows how a node's mapped position on a unit disk can be obtained and also shows related operations in presenting a link feature. Fig. 13 shows operations in presenting a node feature.

The technique in Fig. 11 maps from the hyperbolic plane layout space described above to the unit disk, and from the unit disk to a circular display region. There are two canonical ways of mapping the hyperbolic plane to the unit disk. In both of them, one vicinity in the hyperbolic plane is in focus at the center of the disk while the rest of the hyperbolic plane fades off in a perspective-like fashion toward the perimeter of the disk. The projective mapping, or Klein model, takes lines in the hyperbolic plane to lines in the unit disk, but distorts angles. The conformal mapping, or Poincaré model, preserves angles, but distorts lines in the hyperbolic space into arcs on the unit disk.

It is easy to convert between the two mappings by recomputing the distance from the origin according to $r_k = ((2r_p) / (1 + r_p^2))$, and both have been tested experimentally for use in implementing the browser. The Poincaré model works much better, because it preserves the shapes of fan-outs at nodes and does a better job of using the screen real-estate. In addition, the arcs in the display give a hint of the kind of distortion that is occurring in mapping to the display. Therefore, the technique in Fig. 11 uses the Poincaré model. Rigid transformations of the hyperbolic plane correspond, under the Poincaré model, to linear fractional transformations on the complex plane that send the unit disk to itself. Ahlfors, L.V., *Complex Analysis*. McGraw-Hill, 1966, pp. 76-89, describes the theory of linear fractional transformations.

The Poincaré projection maps circles on the hyperbolic plane into circles on the Euclidean unit disk, though circles will shrink in size the further they are from the origin. The technique of Fig. 11 exploits this property by calculating a circle in the hyperbolic plane around each node that is guaranteed not to intersect the circle of any

other node. When those circles are mapped onto the unit disk they provide a circular display region for each node in which to display a graphical or textual representation of the node. This can be combined with a facility that uses 5 different representations for nodes depending on the amount of real space they receive.

In Fig. 11, the act in box 500 receives a call for mapping with the handle of the root node of a node-link structure, similarly to the act in box 350 in Fig. 9. The act in box 500 also receives a transformation, referred to as a current translator; a function for presenting link features, referred to as link function; a function for presenting node features, referred to as node function; and a limit value that is used to determine whether a node is too near to 10 an edge for presentation of a node feature. For example, the limit value can be 0.07 if the call is made during a sequence of animated images, and can be the size of a pixel otherwise, calculated by taking the reciprocal of the radius of a circular display region to which mapping is 15 being performed. If necessary for satisfactory animation, the limit can be adjusted to reduce the time necessary to render a presentation.

The act in box 502 begins mapping by initializing variables, including a variable to indicate a previous node 20 feature's position and a variable to indicate the previous position's radial gap from a unit disk's perimeter. The act in box 502 can also precompute values to be used within a sequence of animated representations, such as a value for limit as discussed below.

The act in box 504 begins a recursive operation, referred to in Fig. 11 as DoNode, by obtaining the current node's mapped position (x, y) in the unit disk, using techniques described below. The mapped position is the current node's inverse image in the layout space. The act in box 506 then uses the current transformation to present a link feature which has been implemented as a line connecting the position from box 504 to the previous node feature's position.

The act in box 510 determines whether the position 40 from box 504 is such that a node feature can be presented. This has been implemented by calculating a radial gap for the position from box 504, then comparing it with the precomputed limit from box 502 to determine whether it is a position too close to the unit disk's perimeter. To 45 reduce computation within DoNode, the radial gap is calculated as $(1.0 - (x^2 + y^2))$, because the value $(x^2 + y^2)$ must also be computed for another purpose. The act in box 502 can use the limit from box 500 to obtain the precomputed limit, $(1.0 - (1.0 - \text{limit})^2)$. If the radial gap 50 is less than the precomputed limit, the position from box 504 is too close to the unit disk's perimeter.

If the position from box 504 is too close to the perimeter the act in box 510 also compares the radial gap with 55 percent of the previous position's radial gap to determine whether the current position is in a direction from its parent such that its children may be further from the unit disk's perimeter than the current position is. If the act in box 510 finds that the position and direction

are both not okay. DoNode ends. But if the current position is not too close or if its children may be further from the perimeter than it is, a node feature representing the current node may be presented, so the act in box 512 begins an iterative loop that handles each of its children.

In preparation for a recursive call to DoNode, the act in box 514 begins each iteration by setting the previous node feature's position to the position from box 504 and by setting the previous position's radial gap to the radial gap calculated in box 510. These values are set locally within the iteration. Then, the act in box 516 makes a call to DoNode for the next child with the child's handle and with the parent's position from box 504. The act in box 518 returns from DoNode.

When all the children have been handled, so that each has a mapped position in the unit disk, the act in box 520 presents the current node's node feature before DoNode ends. Therefore, the node feature is painted over the link features presented in box 506. The node feature can be positioned by using the radius and mapped position of the node in the layout space to obtain a mapped region in the display plane, with the mapped region defined by a center and a radius as described below. If a wafers variable is on, the node feature includes a circle with the radius and text centered at the center, but if a text variable is on, the node feature includes a rectangle and text centered at the center.

The acts in boxes 506 and 520 each include presenting a feature on display 208. These acts can be implemented as appropriate for the specific display system being used, using conventional display driver techniques. In general, these acts make a constant transform from unit disk positions to display positions.

Fig. 12 shows how the acts in boxes 504 and 506 can be implemented.

The act in box 540 receives a DoNode call with a node's handle and a parent's transformed position. If the node is the root node, the parent's transformed position is taken as its own transformed position. Otherwise, the parent's transformed position has previously been obtained by applying the current translator as described below.

The act in box 542 uses the node's handle to obtain its position in the layout space, which can be obtained from the node's data structure as described in relation to box 372 in Fig. 9. The act in box 544 then applies the current translator from box 500 to the node's position from box 542, obtaining a transformed position in the layout space. The translator can be applied as described above in relation to box 400 in Fig. 10.

The link function from box 500 can be called with the parent's position from box 540 and with the transformed position from box 544 to perform the remaining acts in Fig. 12. The link function begins in box 546 by using the parent's transformed position and the current node's transformed position to obtain mapped positions. In a current implementation, with mapping to a circular display region of radius R_d and centered at (R_d, R_d) from

the upper left corner of a window, the mapping for each coordinate c can be obtained by calculating a floor under $(cR_d + R_d + 0.5)$. This mapping includes an implicit mapping from the layout space to a unit disk, which is done by taking the coordinates of the position in the layout space as x - and y -coordinates of a position on the unit disk. The unit disk position is then mapped to a display position with the above calculation.

The link function continues in box 550 by determining

whether it is appropriate to display a link feature that is an arc. An arc may not be appropriate, for example, during an animation sequence, because an arc requires additional computation. Or an arc may not be appropriate if the difference between the mapped positions from box 546 is small, so that an arc will appear approximately as a line anyway. If an arc is not appropriate, the act in box 552 draws a straight line connecting the mapped positions from box 546.

If an arc is appropriate, the act in box 554 begins by

calculating the arc's radius of curvature in the display space, using the parent's position (x_p, y_p) from box 540 and the transformed position (x_t, y_t) from box 544. The act in box 554 can be implemented by obtaining a center of curvature of the arc that passes through the two points and intersects the boundary circle at right angles. This is the arc that would be mapped from the line between the two points in the layout space. The calculations can then obtain radius of curvature.

The act in box 554 can first obtain the value $d = 2$

$(x_p y_t - x_t y_p)$, then the two values $a = ((1 + x_t^2 + y_t^2)/d)$ and $b = ((1 + x_p^2 + y_p^2)/d)$. The center of curvature is the point (x_c, y_c) such that $x_c = by_t - ay_p$, $y_c = ax_p - bx_t$. The radius of curvature can be calculated as the distance from the center of curvature to either the parent's transformed position (x_p, y_p) or the current node's transformed position (x_t, y_t) .

The act in box 556 determines whether the radius of

curvature from box 554 is too great for accurate drawing. If so, the act in box 552 draws a straight line, as described above. But if the radius is small enough, the act in box 558 uses conventional calculations to draw the arc as calculated in box 554 on the display.

After the act in box 552 or the act in box 558, the technique of Fig. 12 returns to box 510 in Fig. 11.

Fig. 13 shows how the act in box 520 in Fig. 11 can be implemented.

The act in box 570 in Fig. 13 begins by testing whether it is appropriate to present a node feature. This act can include determining whether a parameter is set to indicate that node features should be displayed or whether, for an animation sequence, whether a parameter is set to indicate that node features should not be displayed during animation. If presentation of a node feature is not appropriate, DoNode ends.

If presentation of a node feature is appropriate, the act in box 572 uses the position and radius of the current node from the data structure described in relation to box 372 to obtain a center position and a unit disk radius in

the unit disk. The x - and y -coordinates of the center position can be obtained by multiplying each by $((1.0 - r_f^2)/(1.0 - (x^2 + y^2)r_f^2))$, where r_f is the radius of the current node, optionally adjusted to add a gap between nodes by multiplying by a number less than one such as 0.9. The act in box 572 can multiply the radius r_f by $((1.0 - (x^2 + y^2))/(1.0 - (x^2 + y^2)r_f^2))$ to obtain the unitdisk radius r_u .

The act in box 574 then uses the center position and unit disk radius r_u from box 572 to obtain a display position at the upper left corner of a bounding box around the corresponding circle on the display and to obtain a diameter of the corresponding circle on the display. The x - and y -coordinates of the upper left corner can each be obtained by using the coordinate c to calculate a floor under $(R_d(c - r_u) + R_d + 0.5)$. The diameter can be obtained as the difference between a floor under the value $(R_d(x + r_u) + R_d + 0.5)$ and a floor under the value $(R_d(x - r_u) + R_d + 0.5)$.

The act in box 576 then obtains a count of how many characters fit into the size from box 574. The count can be obtained with a floor function on the size and a character size. The act in box 580 branches on whether the count is greater than two; if not, DoNode ends because there is insufficient size to present a meaningful text.

If the count is greater than two, the act in box 582 obtains a position for the text. For example, an x -coordinate can be the sum of the mapped x -coordinate and half the difference between the mapped size and the product of the character size with the number of characters in the text for the node. The y -coordinate can be a floor under $(R_d y + R_d + 0.5)$, where y is the center y -coordinate from box 572.

The act in box 590 then branches based on whether a wafer or a rectangle should be presented. If a wafer, the act in box 592 draws a circle at the mapped position from box 574. If a rectangle, the act in box 594 draws one that fits around the text at the position from box 582. Then, the act in box 596 draws the text at the position from box 582, before DoNode ends.

As described below, the techniques in Figs. 11-13 can also be used in responding to a request for a change in position.

Figs. 14-16 illustrate operations in responding to a signal from a user requesting a change in position of a representation. Fig. 14 shows how CPU 202 can respond to a signal from a user by mapping to layout space and obtaining and applying a transformation: the technique of Fig. 14 can be used to implement the acts in boxes 272, 274, and 276 in Fig. 7, and similar techniques could be used to implement the acts in boxes 292 and 294 in Fig. 8. Fig. 15 illustrates a way of obtaining a transformation from a first position to a second while preserving orientation of a root node. Fig. 16 illustrates a way of obtaining a transformation from a first position to a second while orienting children of a node at the center of a region of greater spacings to the right.

The technique of Fig. 14 translates the structure on the hyperbolic plane to provide a continuous, geometric

way of moving focus. The user can easily browse the structure while maintaining the visual context of the entire structure. The user can change the focus either by clicking on any visible point to bring it into focus at the center, or by dragging any visible point interactively to any other position. In either case, the rest of the representation transforms appropriately. Regions that approach the center become magnified, while regions that were in the center shrink as they are moved toward the edge.

To compute the composition of two transformations, the point (0,0) can be transformed under the first transformation, and then the second transformation can be applied to the result. The same can then be done to the point (1,0). The images of those two points are enough to identify the composition, and the appropriate geometric operations can be computed straightforwardly from that information.

The act in box 610 in Fig. 14 receives a call for a change in position indicating a first position and a second position in the display region in which a representation is being presented. In a dragging operation, the first and second positions can be explicitly indicated. In a click operation, the first position can be the position of the click, and the second position can be the center of the display region.

The act in box 612 begins by converting the first and second positions from box 610 to unit disk coordinates to which hyperbolic geometry operations can be applied. In a current implementation, a representation is presented in a circular display region of radius r within a window, so that the act in box 612 can be implemented by dividing each coordinate by r , then subtracting 1.0.

The act in box 620 then tests whether either of the unit disk positions from box 612 is too close to the edge of the unit disk, e.g. within approximately 0.025 of the edge. If so, the act in box 622 adjusts the position that is too close to the edge, such as by multiplying each coordinate by the ratio of 0.97 to the distance from the origin to the position, thus slightly reducing each coordinate to move the position further from the edge.

The conversion in box 612 and, if necessary, the adjustment in box 622 produce an inverse image of the first position and an inverse image of the second position, both in the layout space.

The act in box 624 then obtains a translator, by finding a transformation in the layout space that will take the inverse image of the first position to the inverse image of the second position. The translator can be obtained using techniques described in more detail below. The technique in Fig. 14 can follow two different paths, depending on whether animation is being performed or not, as illustrated by the act in box 630. In a dragging operation, animation may be unnecessary to produce a perception of continuity, but in a click operation, it may be necessary to break the change in position into an animated sequence of smaller steps that produce the perception of continuity.

If no animation is being performed, the act in box 632 obtains a new current translation by composing the translator from box 624 with the current translation, meaning the translation used to obtain the displayed representation from the layout data. The composition of a first translator with a second translator, each defined by a point P_i and an angle of rotation θ_i where i indicates the translator, can be obtained by calculating the new values P_n and θ_n as follows: $P_n = ((P_1 + P_2\theta_1)/(1.0 + P_1^*\theta_1))$, where P_1^* is the complex conjugate of P_1 , and $\theta_n = ((\theta_2(\theta_1 + P_1P_2^*)/(1.0 + P_1^*\theta_1))$. Composition does not commute, so the order in which translators are composed affects the resulting translation.

The act in box 634 applies the new current translation from box 632 to the layout data to obtain transformed positions, then maps the transformed positions and presents a representation as described above in relation to Fig. 11. Then the response to the call in box 610 ends.

If animation is being performed, the act in box 640 obtains the n th root of the translator from box 624 as described below. Then, the act in box 642 prepares for an iterative loop by initializing an intermediate translation to the current translation used to present the representation. The act in box 644 begins an iterative loop that is performed n times to obtain n intermediate translations. The act in box 646 composes the n th root from box 640 with the intermediate translation to obtain a new intermediate translation, using the composition technique described in relation to box 632. Then, the act in box 648 adds the new intermediate translation from box 646 to a queue of intermediate translations.

After n intermediate translations have been queued, the act in box 650 obtains a final translation by composing the translator from box 624 with the current translation used to present the representation. The act in box 652 adds the final translation to the queue, and the act in box 654 then goes through the queue, applying each translation to the layout data to obtain transformed positions, then mapping the transformed positions and presenting a representation as described above in relation to box 634, but with one representation for each queued translation. Then the response to the call in box 610 ends.

The act in box 640 can obtain the n th root as follows using the point P_i and the angle of rotation θ_i of a translator. In order to determine which of three cases applies, the act in box 640 can obtain $d = (4|P_i|^2 - |\theta_i| - 1)^2$. The act in box 640 may also precompute certain values that are useful in more than one of the cases, such as $q_a = (|\theta_i| - 1)^2/|\theta_i| + 1|^2$ and $q_b = (d/|\theta_i| + 1)^2$.

If $d > 0$, the act in box 640 can obtain $a_1 = (\tanh((\operatorname{atanh}(q_b^{1/2})/n)/(q_b^{1/2})))$; $t_1 = a_1^2 q_a$; and $r_1 = a_1((1 + q_a)/(1 + t_1))^{1/2}$.

If $d = 0$, the act in box 640 can obtain similar values $t_2 = (q_a/n^2)$ and $r_2 = (((1 + q_a)/(1 + t_2)))^{1/2}$.

If $d < 0$, the act in box 640 can obtain values $a_3 = (\tan((\operatorname{atan}(-q_b^{1/2})/n)))$; $t_3 = (a_3^2|\theta_i| - 1)^2/(-d)$; and $r_3 = (4t_3/((1 + t_3)(|\theta_i| - 1)^2))^{1/2}$.

Then, the act in box 640 can obtain the the angle of rotation θ_n of the n th root translator by first obtaining a complex raw angle θ_r , which has a real part $((1 - t_i)/(1 + t_i))$ and an imaginary part $(2t_i^{1/2}/(1 + t_i))$, where t_i is the value for the applicable case as described above, and the branch cut for the complex square root is chosen to give the angle with the smaller absolute rotation. θ_r can then be normalized to have a magnitude of 1 to obtain θ_n .

Finally, the act in box 640 can obtain the point P_n of the n th root translator as $P_n = P_1 r_i(\theta_n/\theta_i)^{1/2}$, where r_i is the value for the applicable case as described above.

Fig. 15 shows how the act in box 624 could be implemented to obtain a translator from a first position (x_1, y_1) to a second position (x_2, y_2) while preserving orientation of a third position (x_3, y_3) . As shown, the acts in Fig. 15 follow either box 620 or box 622 in Fig. 14, and, in the current implementation, precede box 632 if performed in response to a dragging call and precede box 640 if performed in response to a click call. If called in response to a dragging call, the first position is the beginning position of the dragging motion and the second position is the ending position. If called in response to a click call, the first position is the position of the click and the second position is the origin of the unit disk. In either case, the third position is the current position of the root node, obtained by performing the current translation on the position $(0, 0)$ in the layout space.

The act in box 670 obtains a first translator from the third position to the origin, and can be implemented as described above in relation to box 400 in Fig. 10. The act in box 672 then applies the first translator to the first position to obtain a transformed first position.

The act in box 674 obtains a second translator that has no rotation component and that takes the transformed first position (x_a, y_a) to the second position (x_b, y_b) . The act in box 674 can be implemented by obtaining a translator with point P_2 and an angle of rotation $(1.0, 0.0)$. The act in box 674 can obtain the real part of P_2 by obtaining the real part of the complex value $((B - A)(1 + (AB)^*)/(1 - (IAB)^*)^2$. The act in box 674 can similarly obtain the imaginary part of P_2 by obtaining the imaginary part of the complex value $((B - A)(1 - (AB)^*)/(1 - (IAB)^*)^2$. In these equations, $A = x_a + iy_a$ and $B = x_b + iy_b$.

The act in box 676 then obtains a translator that is the composition of the second translator from box 674 and the first translator from box 670. The translator from box 676 can then be used in the act in box 632 or in the acts in boxes 640 and 650 in Fig. 14.

Fig. 16 shows how the act in box 624 could be implemented to obtain a translator from a first position to a second position that is the origin, with the children of the node feature nearest the center of the region of greater spacings oriented to right. As shown, the acts in Fig. 16 follow either box 620 or box 622 in Fig. 14, and, in the current implementation, precede box 640 because they are only performed in response to a click call. The first position is the position of the click.

The act in box 690 obtains a first translator from the

first position to the origin, and can be implemented as described above in relation to box 400 in Fig. 10. The act in box 692 obtains a second translator that is the composition of the first translator from box 690 and the current translation used to obtain the displayed representation.

The act in box 694 applies the inverted current transformation to obtain an inverted image of the first position. A transformation defined by a point P and an angle of rotation θ can be inverted by obtaining the transformation with the point $(-P\theta^*)$ and the angle of rotation θ^* . The act in box 694 also uses the inverted image of the first position to find the parent of the nearest node.

The act in box 700 applies the second translator from box 692 to the position of the parent node from box 694 and normalizes the result to have a magnitude of one to obtain a parent direction (θ_x, θ_y) . The act in box 702 then obtains a third translator that rotates the layout space to the parent direction without moving the origin. The third translator can have the point $(0, 0)$ and the direction $(-\theta_x, \theta_y)$.

The act in box 704 then obtains a fourth translator by taking the composition of the third translator from box 702 with the second translator from box 692. Then the act in box 706 obtains the translator by taking the composition of the fourth translator from box 704 and the inverted current translation. The translator from box 706 can then be used in the acts in boxes 640 and 650 in Fig. 14.

The implementations described above operate on node-link data defining an organization chart, but could operate on node-link data defining any appropriate node-link structure. For example, the node-link data could define all web pages of a network within some limits, and could be obtained by converting the web pages to a tree using an appropriate pruning constraint.

The implementations described above present a representation in a circular display region, but a representation could be presented in any other appropriately shaped display region, such as an ellipse, through an appropriate mapping from the hyperbolic plane.

The implementations described above present representations of an entire node-link structure, but mappings could be performed that permit fly-through navigation by presenting partial representations of a structure.

The implementations described above show representations with a single region of greater spacings, but representations could be presented with two or more regions of greater spacings, to focus on more node features. Of particular interest is a dual focus representation in which the points $(a, 0)$ and $(-a, 0)$ in the Poincaré map become the centers of binocular-like regions of greater spacings. For this representation, a point z in the hyperbolic plane represented by a Poincaré map to $((z(1 - a^2) / (1 - a^2 z^2))$ in the display space.

The implementations described above show node features that are rectangular or circular and link features that are lines such as arcs or straight lines, but represen-

tations could be presented with other types of node and link features. For example, greater use of display space may be obtained with oval or oblong node features that extend into otherwise unused display areas. Furthermore, the region assigned to a node in a hyperbolic plane or other layout space could have a non-circular shape.

The implementations described above employ a hyperbolic plane, but another space with negative curvature, such as a hyperbolic n-space, could be used as the layout space.

The implementations described above use complex numbers to represent positions in the hyperbolic plane, in part because continuous functions in complex numbers are conformal mappings. Positions in the hyperbolic plane could, however, be represented with other types of data or coordinates.

The implementations described above map from the hyperbolic plane to a unit disk and from the unit disk to a circular display region, but other types of mapping from a layout space to a display region could be used. For example, a mapping from the Poincaré model to the Klein model could be used, although such a mapping has not yet been satisfactorily performed in experiments.

In performing animation, the implementations described above use n steps, where n is a constant. The number of animation steps could instead be a variable that depends on the distance between the starting position and the ending position across which animated motion is occurring.

The implementation described above in relation to Figs. 9-16 implements the technique in Fig. 7. The invention could be implemented according to the technique in Fig. 8, however, or in other ways to provide a browser interface.

The invention has been applied in providing an interactive browser as described in concurrently filed, commonly assigned application EP-A-000,000, corresponding to U.S. Patent Application No. 08/306,074, entitled "Displaying Node-Link Structure with Region of Greater Spacings and Peripheral Branches". The invention could be applied in many other ways, including various other interactive browser applications.

In addition uniform layout of a node-link structure in a space with negative curvature may be applicable to interactive voice systems, and mappings may be performed from the layout space to sound renderings.

Claims

50

1. A method comprising:

obtaining layout data: the layout data indicating positions in a layout space for parts of a node-link structure: the layout space being a space with negative curvature: the node-link structure including nodes and links, each link relating at least two of the nodes: the layout data indicating positions in the layout space for a set of the nodes: the set of nodes

- forming a branch that includes two or more levels of nodes including a top level and at least one lower level, the top level including a top level node and the lower levels including lower level nodes, each node at each lower level having a parent node at a next higher level to which the node is related through one link: for each of a set of two or more parent nodes, the layout data indicating:
- 5 a parent position in the layout space for the parent node: and
 - 10 child positions that lie approximately along a circle in the layout space for a number of lower level nodes that share the parent node, the parent position being approximately at the center of the circle: the number of lower level nodes being two or more: adjacent child positions along the circle being separated by approximately a base spacing;
 - 15 the circles having radii that, together, approximate a function that increases slowly with number of lower level nodes such that the radii and spacings between adjacent child positions along circles are all approximately uniform within the branch; and
 - 20 using the layout data to present a first representation of the node-link structure on a display:
 - 25 receiving a signal from a user: the signal indicating a change from a first display position to a second display position: the first representation including a first feature near the first display position, the first feature representing a first part of the node-link structure: and
 - 30 in response to the signal, presenting a second representation of the node-link structure on the display: the second representation including a second feature representing the first part of the node-link structure, the second feature being presented near the second display position: the second representation being perceptible as a changed continuation of the first representation.
2. The method of claim 1 in which the space with negative curvature is a discrete approximation of a hyperbolic plane.
3. A method of operating a machine: the machine including:
- 35 memory;
 - 40 user input circuitry for providing data indicating signals from a user;
 - 45 a display;
 - 50 a processor connected for accessing data stored in the memory, for receiving data from the user input circuitry, and for providing data to cause presentation of images by the display: and
 - 55 layout data stored in the memory: the layout data indicating positions in a layout space for parts of a node-link structure: the layout space being a space with negative curvature: the node-link structure including nodes and links, each link relating at

- least two of the nodes: the layout data indicating positions in the layout space for a set of the nodes: the set of nodes forming a branch that includes two or more levels of nodes including a top level and at least one lower level, the top level including a top level node and the lower levels including lower level nodes, each node at each lower level having a parent node at a next higher level to which the node is related through one link: for each of a set of two or more parent nodes, the layout data indicating:
- 5 a parent position in the layout space for the parent node: and
 - 10 child positions that lie approximately along a circle in the layout space for a number of lower level nodes that share the parent node, the parent position being approximately at the center of the circle: the number of lower level nodes being two or more: adjacent child positions along the circle being separated by approximately a base spacing;
 - 15 the circles having radii that, together, approximate a function that increases slowly with number of lower level nodes such that the radii and spacings between adjacent child positions along circles are all approximately uniform within the branch:
 - 20 the method comprising:
 - 25 operating the processor to use the layout data to provide first representation data to the display, the first representation data causing the display to present a first representation of the node-link structure:
 - 30 receiving a signal from a user: the signal indicating a change from a first display position to a second display position: the first representation including a first feature near the first display position, the first feature representing a first part of the node-link structure: and
 - 35 in response to the signal, operating the processor to automatically provide second representation data to the display, the second representation data causing the display to present a second representation of the node-link structure on the display: the second representation including a second feature representing the first part of the node-link structure, the second feature being presented near the second display position: the second representation being perceptible as a changed continuation of the first representation.
 - 40
 - 45
 - 50
 - 55
4. A machine comprising:
- memory;
 - user input circuitry for providing data indicating signals from a user;
 - a display;
 - a processor connected for accessing data stored in the memory, for receiving data from the user input circuitry, and for providing data to cause presentation of images by the display:
 - layout data stored in the memory: the layout

- data indicating positions in a layout space for parts of a node-link structure: the layout space being a space with negative curvature: the node-link structure including nodes and links, each link relating at least two of the nodes: the layout data indicating positions in the layout space for a set of the nodes; the set of nodes forming a branch that includes two or more levels of nodes including a top level and at least one lower level, the top level including a top level node and the lower levels including lower level nodes, each node at each lower level having a parent node at a next higher level to which the node is related through one link: for each of a set of two or more parent nodes, the layout data indicating:
- 5 a parent position in the layout space for the parent node: and
 - 10 child positions that lie approximately along a circle in the layout space for a number of lower level nodes that share the parent node, the parent position being approximately at the center of the circle: the number of lower level nodes being two or more: adjacent child positions along the circle being separated by approximately a base spacing;
 - 15 the circles having radii that, together, approximate a function that increases slowly with number of lower level nodes such that the radii and spacings between adjacent child positions along circles are all approximately uniform within the branch: and
 - 20 program data stored in the memory: the program data indicating instructions the processor executes: the processor, in executing the instructions:
 - 25 using the layout data to provide first representation data to the display, the first representation data causing the display to present a first representation of the node-link structure;
 - 30 receiving a signal from a user: the signal indicating a change from a first display position to a second display position: the first representation including a first feature near the first display position, the first feature representing a first part of the node-link structure: and
 - 35 in response to the signal, automatically providing second representation data to the display, the second representation data causing the display to present a second representation of the node-link structure on the display: the second representation including a second feature representing the first part of the node-link structure, the second feature being presented near the second display position: the second representation being perceptible as a changed continuation of the first representation.
5. An article of manufacture for use in a system that includes:
- 40 memory:
 - 45 a storage medium access device:
 - 50 user input circuitry for providing data indicating signals from a user:

a display:
 a processor connected for accessing data stored in the memory, for receiving data accessed on a storage medium by the storage medium access device, and for providing data to cause presentation of images by the display: and
 layout data stored in the memory: the layout data indicating positions in a layout space for parts of a node-link structure: the layout space being a space with negative curvature: the node-link structure including nodes and links, each link relating at least two of the nodes: the layout data indicating positions in the layout space for a set of the nodes; the set of nodes forming a branch that includes two or more levels of nodes including a top level and at least one lower level, the top level including a top level node and the lower levels including lower level nodes, each node at each lower level having a parent node at a next higher level to which the node is related through one link: for each of a set of two or more parent nodes, the layout data indicating:

- 55 a parent position in the layout space for the parent node: and
- 60 child positions that lie approximately along a circle in the layout space for a number of lower level nodes that share the parent node, the parent position being approximately at the center of the circle: the number of lower level nodes being two or more: adjacent child positions along the circle being separated by approximately a base spacing;
- 65 the circles having radii that, together, approximate a function that increases slowly with number of lower level nodes such that the radii and spacings between adjacent child positions along circles are all approximately uniform within the branch:
- 70 the article of manufacture comprising:
- 75 a storage medium: and
- 80 instruction data stored by the storage medium: the instruction data including data units positioned on the storage medium so that the storage medium access device can access the data units and provide the data units to the processor in a sequence: the data units, when in the sequence, indicating instructions the processor can execute: the processor, in executing the instructions:
- 85 using the layout data to provide first representation data to the display, the first representation data causing the display to present a first representation of the node-link structure;
- 90 receiving a signal from a user: the signal indicating a change from a first display position to a second display position: the first representation including a first feature near the first display position, the first feature representing a first part of the node-link structure: and
- 95 in response to the signal, automatically providing second representation data to the display, the second representation data causing the display to

- present a second representation of the node-link structure on the display: the second representation including a second feature representing the first part of the node-link structure, the second feature being presented near the second display position: the second representation being perceptible as a changed continuation of the first representation.
6. A method of operating a first machine that includes:
 5 first memory;
 a first connection to a network;
 a first processor connected for accessing data stored in the first memory and for establishing connections with other machines over the network through the first connection: and
 instruction data stored in the first memory: the instruction data indicating instructions that can be executed by a second machine, the second machine including:
 10 second memory;
 a second connection to the network;
 a second processor connected for accessing data stored in the second memory and for establishing connections with other machines over the network through the second connection: and
 layout data stored in the memory: the layout data indicating positions in a layout space for parts of a node-link structure: the layout space being a space with negative curvature: the node-link structure including nodes and links, each link relating at least two of the nodes; the layout data indicating positions in the layout space for a set of the nodes: the set of nodes forming a branch that includes two or more levels of nodes including a top level and at least one lower level, the top level including a top level node and the lower levels including lower level nodes, each node at each lower level having a parent node at a next higher level to which the node is related through one link: for each of a set of two or more parent nodes, the layout data indicating:
 15 a parent position in the layout space for the parent node: and
 child positions that lie approximately along a circle in the layout space for a number of lower level nodes that share the parent node, the parent position being approximately at the center of the circle: the number of lower level nodes being two or more: adjacent child positions along the circle being separated by approximately a base spacing:
 20 the circles having radii that, together, approximate a function that increases slowly with number of lower level nodes such that the radii and spacings between adjacent child positions along circles are all approximately uniform within the branch:
 25 the method comprising:
 30 operating the first processor to establish a connection to the second processor over the network through the first and second connections: and
 35 operating the first processor to transfer the instruction data over the network to the second processor so that the second processor can store the instruction data in the second memory and can execute the instructions indicated by the instruction data: the second processor, in executing the instructions indicated by the instruction data:
 40 using the layout data to provide first representation data to the display, the first representation data causing the display to present a first representation of the node-link structure:
 45 receiving a signal from a user: the signal indicating a change from a first display position to a second display position: the first representation including a first feature near the first display position, the first feature representing a first part of the node-link structure: and
 50 in response to the signal, automatically providing second representation data to the display, the second representation data causing the display to present a second representation of the node-link structure on the display: the second representation including a second feature representing the first part of the node-link structure, the second feature being presented near the second display position: the second representation being perceptible as a changed continuation of the first representation.
7. A method of operating a machine: the machine including:
 5 memory:
 a processor connected for accessing data stored in the memory: and
 node-link data stored in the memory: the node-link data defining a node-link structure: the node-link structure including nodes and links, each link relating at least two of the nodes;
 10 the method comprising:
 15 operating the processor to use the node-link data to obtain layout data: the layout data indicating positions in a layout space for parts of a node-link structure: the layout space being a space with negative curvature: the node-link structure including nodes and links, each link relating at least two of the nodes: the layout data indicating positions in the layout space for a set of the nodes: the set of nodes forming a branch that includes two or more levels of nodes including a top level and at least one lower level, the top level including a top level node and the lower levels including lower level nodes, each node at each lower level having a parent node at a next higher level to which the node is related through one link: for each of a set of two or more parent nodes, the layout data indicating:
 20 a parent position in the layout space for the parent node: and
 25 child positions that lie approximately along a circle in the layout space for a number of lower level nodes that share the parent node, the parent position being approximately at the center of the circle: the number of lower level nodes being two or more: adjacent child positions along the circle being separated by approximately a base spacing:
 30 the circles having radii that, together, approximate a function that increases slowly with number of lower level nodes such that the radii and spacings between adjacent child positions along circles are all approximately uniform within the branch:
 35 the method comprising:
 40 operating the first processor to establish a connection to the second processor over the network through the first and second connections: and
 45 operating the first processor to transfer the instruction data over the network to the second processor so that the second processor can store the instruction data in the second memory and can execute the instructions indicated by the instruction data: the second processor, in executing the instructions indicated by the instruction data:
 50 using the layout data to provide first representation data to the display, the first representation data causing the display to present a first representation of the node-link structure:
 55 receiving a signal from a user: the signal indicating a change from a first display position to a second display position: the first representation including a first feature near the first display position, the first feature representing a first part of the node-link structure: and
 60 in response to the signal, automatically providing second representation data to the display, the second representation data causing the display to present a second representation of the node-link structure on the display: the second representation including a second feature representing the first part of the node-link structure, the second feature being presented near the second display position: the second representation being perceptible as a changed continuation of the first representation.

nodes that share the parent node, the parent position being approximately at the center of the circle: the number of lower level nodes being two or more: adjacent child positions along the circle being separated by approximately a base spacing;

the circles having radii that, together, approximate a function that increases slowly with number of lower level nodes such that the radii and spacings between adjacent child positions along circles are all approximately uniform within the branch; and

operating the processor to use the layout data to perform two or more different mappings: each mapping obtaining mapped data: the mapped data obtained by each mapping indicating positions for a subset of the set of nodes.

8. A machine comprising:

memory;

a display;

a processor connected for accessing data stored in the memory and for providing data to cause presentation of images by the display;

node-link data stored in the memory: the node-link data defining a node-link structure that includes nodes and links, each link relating at least two of the nodes: and

program data stored in the memory: the data program indicating instructions the processor executes: the processor, in executing the instructions:

using the node-link data to obtain layout data: the layout data indicating positions in a layout space for parts of a node-link structure: the layout space being a space with negative curvature: the node-link structure including nodes and links, each link relating at least two of the nodes: the layout data indicating positions in the layout space for a set of the nodes: the set of nodes forming a branch that includes two or more levels of nodes including a top level and at least one lower level, the top level including a top level node and the lower levels including lower level nodes, each node at each lower level having a parent node at a next higher level to which the node is related through one link: for each of a set of two or more parent nodes, the layout data indicating:

a parent position in the layout space for the parent node: and

child positions that lie approximately along a circle in the layout space for a number of lower level nodes that share the parent node, the parent position being approximately at the center of the circle: the number of lower level nodes being two or more: adjacent child positions along the circle being separated by approximately a base spacing;

the circles having radii that, together, approximate a function that increases slowly with number of lower level nodes such that the radii and spacings between adjacent child positions along circles are

all approximately uniform within the branch; and using the layout data to perform two or more different mappings: each mapping obtaining mapped data: the mapped data obtained by each mapping indicating positions for a subset of the set of nodes.

9. An article of manufacture for use in a machine that includes:

memory;

a storage medium access device;

a display;

a processor connected for accessing data stored in the memory, for receiving data accessed on a storage medium by the storage medium access device, and for providing data to cause presentation of images by the display: and

node-link data stored in the memory: the node-link data defining a node-link structure that includes nodes and links, each link relating at least two of the nodes:

the article of manufacture comprising:

a storage medium: and

instruction data stored by the storage medium: the instruction data including data units positioned on the storage medium so that the storage medium access device can access the data units and provide the data units to the processor in a sequence: the data units, when in the sequence, indicating instructions the processor can execute: the processor, in executing the instructions:

using the node-link data to obtain layout data: the layout data indicating positions in a layout space for parts of a node-link structure: the layout space being a space with negative curvature: the node-link structure including nodes and links, each link relating at least two of the nodes: the layout data indicating positions in the layout space for a set of the nodes: the set of nodes forming a branch that includes two or more levels of nodes including a top level and at least one lower level, the top level including a top level node and the lower levels including lower level nodes, each node at each lower level having a parent node at a next higher level to which the node is related through one link: for each of a set of two or more parent nodes, the layout data indicating:

a parent position in the layout space for the parent node: and

child positions that lie approximately along a circle in the layout space for a number of lower level nodes that share the parent node, the parent position being approximately at the center of the circle: the number of lower level nodes being two or more: adjacent child positions along the circle being separated by approximately a base spacing;

the circles having radii that, together, approximate a function that increases slowly with number

- of lower level nodes such that the radii and spacings between adjacent child positions along circles are all approximately uniform within the branch; and
- using the layout data to perform two or more different mappings: each mapping obtaining mapped data: the mapped data obtained by each mapping indicating positions for a subset of the set of nodes.
10. A method of operating a first machine that includes:
- first memory;
 - a first connection to a network;
 - a first processor connected for accessing data stored in the first memory and for establishing connections with other machines over the network through the first connection; and
 - instruction data stored in the first memory: the instruction data indicating instructions that can be executed by a second machine, the second machine including:
 - second memory;
 - a second connection to the network;
 - a second processor connected for accessing data stored in the second memory and for establishing connections with other machines over the network through the second connection; and
 - node-link data stored in the second memory: the node-link data defining a node-link structure that includes nodes and links, each link relating at least two of the nodes;
 - the method comprising:
 - operating the first processor to establish a connection to the second processor over the network through the first and second connections; and
 - operating the first processor to transfer the instruction data over the network to the second processor so that the second processor can store the instruction data in the second memory and can execute the instructions indicated by the instruction data: the second processor, in executing the instructions indicated by the instruction data:
 - using the node-link data to obtain layout data: the layout data indicating positions in a layout space for parts of a node-link structure: the layout space being a space with negative curvature: the node-link structure including nodes and links, each link relating at least two of the nodes: the layout data indicating positions in the layout space for a set of the nodes: the set of nodes forming a branch that includes two or more levels of nodes including a top level and at least one lower level, the top level including a top level node and the lower levels including lower level nodes, each node at each lower level having a parent node at a next higher level to which the node is related through one link: for each of a set of two or more parent nodes, the layout data indicating:
 - a parent position in the layout space for the

parent node: and

child positions that lie approximately along a circle in the layout space for a number of lower level nodes that share the parent node, the parent position being approximately at the center of the circle; the number of lower level nodes being two or more: adjacent child positions along the circle being separated by approximately a base spacing:

the circles having radii that, together, approximate a function that increases slowly with number of lower level nodes such that the radii and spacings between adjacent child positions along circles are all approximately uniform within the branch; and

using the layout data to perform two or more different mappings: each mapping obtaining mapped data: the mapped data obtained by each mapping indicating positions for a subset of the set of nodes.

20

25

30

35

40

45

50

55




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7





FIG. 8






FIG. 10




FIG. 11




FIG. 12




FIG. 13




FIG. 14




FIG. 15





FIG. 16



(19)

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 702 330 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
10.07.1996 Bulletin 1996/28

(51) Int Cl. 6: G06T 15/00, G06T 11/20

(43) Date of publication A2:
20.03.1996 Bulletin 1996/12

(21) Application number: 95306309.6

(22) Date of filing: 08.09.1995

(84) Designated Contracting States:
DE FR GB

- Rao, Ramana B.
San Francisco, CA 94112 (US)

(30) Priority: 14.09.1994 US 306043

(74) Representative: Reynolds, Julian David et al

Rank Xerox Ltd

Patent Department

Parkway

Marlow Buckinghamshire SL7 1YL (GB)

(72) Inventors:

- Lamping, John O.
Los Altos, CA 94024 (US)

(54) Layout of node-link structure in space with negative curvature

(57) Layout data indicate positions in a negatively curved layout space for nodes in a hierarchical branch of a node-link structure. The layout data indicate a parent position for parent nodes and, for children that share a parent node, child positions approximately along a circle in the layout space with the parent position approximately at the circle's center. Adjacent child positions are separated by approximately a base spacing. The radii of circles within the branch together approximate a function that increases slowly with number of child nodes such that the radii and spacings along circles are all approximately uniform within the branch. The layout data can be obtained (60) from data defining the node-link structure. The layout data can be used to perform mappings (64), each obtaining positions for a subset of the nodes. The layout data can be used to present (122: Fig. 4) a first representation of the node-link structure on a display. In response to a user signal indicating (124: Fig. 4) a change from a first display position near a first feature to a second display position, a second representation can be presented (126: Fig 4) that is perceptible as a changed continuation of the first. The second representation includes, near the second display position, a second feature representing the same part of the node-link structure as the first feature. The second representation can be obtained by a transformation of the layout space, which can be a discrete approximation of a hyperbolic plane.




FIG. 2



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 95 30 6309

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
Y	HUMAN FACTORS IN COMPUTING SYSTEMS. REACHING THROUGH TECHNOLOGY. CHI '91. CONFERENCE PROCEEDINGS, NEW ORLEANS, LA, USA, 27 APRIL-2 MAY 1991, ISBN 0-89791-383-3, 1991, NEW YORK, NY, USA, ACM, USA, pages 189-194, XP002002888 ROBERTSON G G ET AL: "Cone Trees: animated 3D visualizations of hierarchical information" * paragraph : "gardening, manipulation , and search" * * figures 1,5 *	1-5,7-9	G06T15/00 G06T11/20
Y	MATHEMATICS AND COMPUTER EDUCATION, WINTER 1993, USA, vol. 27, no. 1, ISSN 0730-8639, pages 51-68, XP002002889 AUSTIN J D ET AL: "An empirical exploration of the Poincare model for hyperbolic geometry" * page 56 - page 61; figures 3-8 *	1-5,7-9	
A	STRIKING A BALANCE, MONTEREY, MAY 3 - 7, 1992, no. -, 3 May 1992, BAUERSFELD P; BENNETT J; LYNCH G, page 369/370 XP000426820 JOHNSON B: "TREEVIZ: TREEMAP VISUALIZATION OF HIERARCHICALLY STRUCTURED INFORMATION" * the whole document *	1-10	G06T
		-/-	
The present search report has been drawn up for all claims			
Place of search	Date of completion of the search	Examiner	
THE HAGUE	13 May 1996	Perez Molina, E	
CATEGORY OF CITED DOCUMENTS		I : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document	
X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document			



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 95 30 6309

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
A	ADVANCES IN CLASSIFICATION RESEARCH. VOL.3. PROCEEDINGS OF THE 3RD ASIS SIG/CR CLASSIFICATION RESEARCH WORKSHOP, PROCEEDINGS OF 3RD ASIS SIG/CR CLASSIFICATION RESEARCH WORKSHOP, PITTSBURGH, PA, USA, 25 OCT. 1992, ISBN 0-938734-79-2, 1993, MEDFORD, NJ, USA. LEARNED INF, USA, pages 37-46, XP002002890 HIRTEL S C ET AL: "Visualization tools for clustering, trees and ordered trees". ----		
A	SYSTEMS AND COMPUTERS IN JAPAN, 1993, USA, vol. 24, no. 8, ISSN 0882-1666, pages 35-46, XP002002891 MATSUURA T ET AL: "A graph editor for large trees with browsing and zooming capabilities"		
A	WO-A-92 09967 (PARASPECTIVES INC) 11 June 1992 -----		TECHNICAL FIELDS SEARCHED (Int.Cl.6)
The present search report has been drawn up for all claims			
Place of search	Date of completion of the search	Examiner	
THE HAGUE	13 May 1996	Perez Molina, E	
CATEGORY OF CITED DOCUMENTS		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application U : document cited for other reasons ----- K : member of the same patent family, corresponding document	
X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document			

THIS PAGE BLANK (USPTO)